

On the Enumeration of all Minimal Triangulations

Nofar Carmeli, Batya Kenig, and Benny Kimelfeld

Technion – Israel Institute of Technology

Abstract. We present an algorithm that enumerates all the minimal triangulations of a graph in incremental polynomial time. Consequently, we get an algorithm for enumerating all the proper tree decompositions, in incremental polynomial time, where “proper” means that the tree decomposition cannot be improved by removing or splitting a bag.

1 Introduction

Many intractable computational problems on graphs admit tractable algorithms when applied to trees or forests. *Tree decomposition* extracts a tree structure from a graph by grouping nodes into *bags* (where each bag is treated as a single node). The corresponding operation on hypergraphs is that of a *generalized hypertree decomposition* [11], in which the bags are associated with both nodes and hyperedges. A generalized hypertree decomposition entails a tree decomposition of the *primal* graph (which has the same set of nodes, and an edge between every two nodes that co-occur in a hyperedge) and an assignment of hyperedge labels (edge covers) to the tree nodes [12]. Tree decomposition and generalized hypertree decomposition have a plethora of applications, including join optimization in databases [11, 26], containment of database queries, constraint-satisfaction problems [19], computation of Nash equilibria in games [11], analysis of probabilistic graphical models [21], and weighted model counting [17].

Past research has focused on obtaining a “good” tree decomposition for the given graph, where goodness is typically by means of the *width*: the maximal cardinality of a bag (minus one). Nevertheless, finding a tree decomposition of a minimal width is NP-hard [2]. Moreover, in various applications the measure of goodness is different from (though related to) the width [12, 17]. In this paper, we explore the task of *enumerating* all (or a subset of) the tree decompositions of a graph. Such algorithms have been proposed in the past for small graphs (representing database queries), without complexity guarantees [26]. Our main result is an enumeration algorithm that runs in incremental polynomial time.

We first need to define which tree decompositions should be enumerated. For example, if we take a graph that is already a tree, we do not wish to enumerate the tree decompositions that group nodes with no reason; in fact, the tree itself is the only reasonable decomposition in this case. Therefore, we consider only tree decompositions that cannot be “improved” by removing or splitting a bag. Such a tree decomposition, which we formally define in Section 2, is said to be *proper*.

We first show that such tree decompositions are in a bijective correspondence to the *minimal triangulations* of the graph at hand. A *triangulation* is a set of edges that is added to the graph to make it chordal, and it is *minimal* if no strict subset of it is a triangulation of the graph. So, our task is reduced to that of enumerating all of the minimal triangulations of a graph. Our main contribution is an algorithm for enumerating all the minimal triangulations of an input graph in incremental polynomial time.

Our approach is as follows. Parra and Scheffler [22] have shown that there is a one-to-one correspondence between the minimal triangulations and the maximal sets of non-crossing minimal separators. (The precise definitions are in Section 2.) So, enumerating the minimal triangulations boils down to enumerating these maximal sets, which can be thought of as maximal independent sets of the graph \mathcal{G} that represents crossings among minimal separators. It is well known that all the maximal independent sets of a graph can be enumerated with polynomial delay [7, 15]. However, this is insufficient for us, since the graph \mathcal{G} is not given as input, and in fact, can have an exponential number of nodes (in the size of the original given graph). Therefore, we cannot construct this graph ahead of time to establish incremental polynomial time. Instead, we use a result by Berry et al. [3], showing that the minimal separators of a graph can be enumerated with polynomial delay. We devise an algorithm that enumerates the maximal independent sets of a graph by assuming that nodes are given by a polynomial-delay iterator, and by assuming some other complexity bounds that are proved to hold in the case of minimal separators.

The rest of the paper is organized as follows. In Section 2 we give preliminary definitions and notation, recall basic results from the literature, and provide some initial insights. In Section 3 we define the notion of a *succinct graph representation* (where nodes are given via an iterator), and give an algorithm for enumerating the maximal independent sets of such a graph. In Section 4 we prove that the graph of minimal separating sets satisfies the complexity assumptions needed for the enumeration algorithm, and thereby establish our main result. Finally, in Section 5, we give a generic algorithm for extending a set of edges to a minimal triangulation.

2 Preliminaries

2.1 Graphs and Cliques

The graphs in this work are undirected. For a graph g , the set of nodes is denoted by $V(g)$, and the set of edges (where an edge is a pair $\{u, v\}$ of distinct nodes) is denoted by $E(g)$. If $U \subseteq V(g)$, then $g|_U$ denotes the subgraph of g induced by U ; hence, $V(g|_U) = U$ and $E(g|_U) = \{\{u, v\} \in E(g) \mid \{u, v\} \subseteq U\}$. Given a subset S of $V(g)$, we denote by $g \setminus S$ the graph obtained from g by removing all the nodes in S (along with their incident edges), that is, the graph $g|_{V(g) \setminus S}$. The *neighborhood* of a node v of g , denoted $N_g(v)$, is the set $\{u \mid \{u, v\} \in E(g)\}$. The *neighborhood* of a set U of nodes of g , denoted $N_g(U)$, is the set $\cup_{v \in U} N_g(v) \setminus U$;

in words, the neighborhood of U consists of every node that is a neighbor of some node in U , and is not itself in U .

Let g be a graph. Let U be a set of nodes of g . We say that U is a *clique* (of g) if every two nodes of U are connected by an edge. We say that U is a *maximal clique* if U is a clique and U is not strictly contained in any other clique. We denote by $\text{MAXCLQ}(g)$ the set of all the maximal cliques of g . The operation of *saturating* U (in g) is that of connecting every non-adjacent pair of nodes in U by a new edge. Hence, if h is obtained from g by saturating U , then U is a clique of h .

2.2 Minimal Separators

Let g be a graph, and let S be a subset of $V(g)$. Let u and v be two nodes of g . We say that S is a (u, v) -separator of g if u and v belong to distinct connected components of $g \setminus S$. We say that S is a *minimal* (u, v) -separator of g if no strict subset of S is a (u, v) -separator. We say that S is a *minimal separator* of g if there are two nodes u and v such that S is a minimal (u, v) -separator. In each of these forms of a separator, we may omit “of g ” if g is clear from the context. We denote by $\text{MINSEP}(g)$ the set of all the minimal separators of g . We mention that the number of minimal separators (i.e., $|\text{MINSEP}(g)|$) may be exponential in the number of nodes (i.e., $|V(g)|$).

Let g be a graph, and let S and T be two minimal separators of g . We say that S *crosses* T , denoted $S \bowtie_g T$, if there are nodes $v, u \in T$ such that S is a (v, u) -separator. If g is clear from the context, we may omit it and write simply $S \bowtie T$. It is known that \bowtie is a symmetric relation: if S crosses T then T crosses S [18, 22]. Hence, if $S \bowtie T$ then we may also say that S and T are *crossing*. When S and T are non-crossing, then we also say that they are *parallel*.

2.3 Chordality and Triangulation

Let g be a graph. For a cycle c in g , a *chord* of c is an edge $e \in E(g)$ that connects two nodes that are non-adjacent in c . We say that g is *chordal* if every cycle of g of length greater than three has a chord. Dirac [9] has shown a characterization of chordal graphs by means of their minimal separators.

Theorem 1. (Dirac [9]) *A graph g is chordal if and only if every minimal separator of g is a clique.*

Rose [24] has shown that a chordal graph g has fewer minimal separators than nodes (that is, if g is chordal then $|\text{MINSEP}(g)| < |V(g)|$). Moreover, it is known that we can find all of these minimal separators in linear time.

Theorem 2. (Kumar and Madhavan [20]) *There is an algorithm that, given a chordal graph g , computes $\text{MINSEP}(g)$ in $O(|V(g)| + |E(g)|)$ time.*

A *triangulation* of a graph g is a graph h such that $V(g) = V(h)$, $E(g) \subseteq E(h)$, and h is chordal. A *minimal triangulation* of g is triangulation h of g with the

following property: for every graph h' with $V(g) = V(h')$, if $E(g) \subseteq E(h') \subsetneq E(h)$ then h' is non-chordal (or in other words, h' is not a triangulation of g). In particular, if g is already chordal then g is the only minimal triangulation of itself. We denote by $\text{MINTRI}(g)$ the set of all the minimal triangulations of g .

2.4 Tree Decomposition

Let g be a graph. A *tree decomposition* d of g is a pair (t, β) , where t is a tree and $\beta : V(t) \rightarrow 2^{V(g)}$ is a function that maps every node of t into a set of nodes of g , so that all of the following hold.

- Nodes are covered: for every node $u \in V(g)$ there is a node $v \in V(t)$ such that $u \in \beta(v)$.
- Edges are covered: for every edge $e \in E(g)$ there is a node $v \in V(t)$ such that $e \subseteq \beta(v)$.
- For all $u, v, w \in V(t)$, if v is on the path between u and w , then $\beta(v)$ contains $\beta(u) \cap \beta(w)$. This property is termed the *junction-tree* property or the *running-intersection* property.

Let g be a graph, and let $d = (t, \beta)$ be a tree decomposition of g . For a node v of t , the set $\beta(v)$ is called a *bag* of d . We denote by $\text{bags}(d)$ the set $\{\beta(v) \mid v \in V(t)\}$. We denote by $\text{saturate}(g, d)$ the graph obtained from g by saturating (i.e., adding an edge between every pair of nodes in) every bag of d .

Let d_1 and d_2 be two tree decompositions of a graph g . We say that d_1 and d_2 are *bag equivalent*, denoted $d_1 \equiv_b d_2$, if $\text{bags}(d_1) = \text{bags}(d_2)$. We denote by $d_1 \sqsubseteq d_2$ the fact that for every bag $b_1 \in \text{bags}(d_1)$ there exists a bag $b_2 \in \text{bags}(d_2)$ such that $b_1 \subseteq b_2$.

Let g be a graph, and let d and d' be tree decompositions of g . We say that d' *strictly subsumes* d if $d' \sqsubseteq d$ and $\text{bags}(d) \subsetneq \text{bags}(d')$. A tree decomposition is *proper* if it is not strictly subsumed by any tree decomposition, and it is *improper* otherwise. Figure 1 shows examples of proper and improper tree decompositions.

2.5 Enumeration

An *enumeration problem* P is a collection of pairs (x, Y) where x is an *input* and Y is a finite set of *answers* for x , denoted by $P(x)$. A *solver* for an enumeration problem P is an algorithm that, when given an input x , produces a sequence of answers such that every answer in $P(x)$ is printed precisely once.

Johnson, Papadimitriou and Yannakakis [15] introduced several different notions of *efficiency* for enumeration algorithms, and we recall these now. Let P be an enumeration problem, and let A be solver for P . We say that A runs in:

- *polynomial total time* if the total execution time of A is polynomial in $(|x| + |P(x)|)$;
- *polynomial delay* if the time between every two consecutive answers produced is polynomial in $|x|$;
- *incremental polynomial time* if, after generating N answers, the time to generate the next $(N + 1)$ st answer is polynomial in $(|x| + N)$.




Fig. 1. A graph g and tree decompositions d_1 , d_2 and d_3 of g . The decomposition d_1 is proper, but d_2 and d_3 are subsumed by d_1 , and hence, improper.

2.6 Preliminary Insights

In this section we give some preliminary insights on our definitions so far.

The following proposition is a folklore, but we give a proof for completeness.

Proposition 1. *Let d be a tree decomposition of a graph g . Every clique of g is contained in some bag of d .*

Proof. Denote $d = (t, \beta)$ and let C be a clique of g . Every node v in C defines a subtree of t that is induced by the bags that contain v . Since d covers the edges of g , every two nodes in C must share some bag in d , and hence, their subtrees must share a vertex. Heggernes [14] shows that every collection of subtrees of a tree satisfies the *Helly property*: if every two subtrees share a vertex, then there exists a vertex that is shared by all the subtrees. In particular, there exists a vertex in d common to all of these subtrees; this shared node corresponds to a bag that contains C . \square

The following proposition states that in a proper tree decomposition, there is no containment among bags.

Proposition 2. *If d is a proper tree decomposition of a graph g , then $\text{bags}(d)$ is an antichain w.r.t. set inclusion.*

Proof. We need to show that a proper tree decomposition cannot have two bags with one contained in the other. Assume, by way of contradiction, that d is a tree decomposition of g with two bags $B, C \in \text{bags}(d)$ where $B \subseteq C$. Let A be the second bag in the path from B to C . Since d is a tree decomposition and A is on the path from B to C , we get that $B = B \cap C \subseteq A$.

Define d' to be the graph obtained from d by removing B and connecting A to all other neighbors of B , as illustrated in Figure 2. We will show that d' is a tree decomposition for g . The first two properties of the tree decomposition still hold because A contains B . Consider the path between two bags α and β of d' . If the path between them is the same as in d , the third property still holds. If it changed, then the path used to go through B , and the only new bag that may appear in this path is A . In this case, $\alpha \cap \beta \subseteq B \subseteq A$, and the third property holds as well. We have found a tree decomposition d' for g that strictly subsumes d , hence d is improper, hence a contradiction. \square




Fig. 2. Obtaining a strictly subsuming tree decomposition d' given a tree decomposition d with $B \subseteq C$.

Jordan [16] shows the following characterization of chordal graphs by means of tree decompositions.

Theorem 3. (Jordan [16]) *A graph g is chordal if and only if it has a tree decomposition where all the bags are cliques of g .*

From Theorem 3, the following proposition easily follows.

Proposition 3. *If d is a tree decomposition of a graph g , then $\text{satrate}(g, d)$ is a triangulation of g .*

Proof. It is straightforward to show that d is a tree decomposition of $\text{satrate}(g, d)$. Hence, since every bag of d is a clique of $\text{satrate}(g, d)$, it follows from Theorem 3 that $\text{satrate}(g, d)$ is chordal. \square

The next proposition states that a chordal graph g has a single proper tree decomposition, up to the equivalence \equiv_b , with the set of bags being precisely the set of maximal cliques.

Proposition 4. *If d is a proper tree decomposition of a chordal graph g , then $\text{bags}(d) = \text{MAXCLQ}(g)$.*

Proof. According to Proposition 1, every clique of g is contained in some bag of d , and according to Theorem 3, g has some tree decomposition, say d' , where all the bags are cliques of g . So we have that $d' \sqsubseteq d$. If $\text{bags}(d) \not\subseteq \text{bags}(d')$, then d' strictly subsumes d , in contradiction to the fact that d is proper. Hence $\text{bags}(d) \subseteq \text{bags}(d')$, meaning that the bags of d are cliques of g . It thus follows that every maximal clique is a bag of d , or in notation, $\text{MAXCLQ}(g) \subseteq \text{bags}(d)$. Finally, Proposition 2 states that the bags of d are an antichain w.r.t. set inclusion, and hence, $\text{bags}(d) \subseteq \text{MAXCLQ}(g)$. We conclude that $\text{bags}(d) = \text{MAXCLQ}(g)$, as claimed. \square

The next theorem relates proper tree decompositions to minimal triangulations, and reduces the enumeration of the former into the enumeration of the latter.

Theorem 4. *Let g be a graph. There is a bijection M between $\text{MINTRI}(g)$ and the equivalence classes of \equiv_b over the proper tree decompositions of g . Moreover, given a minimal triangulation h of g , the proper tree decompositions in the class $M(h)$ can be enumerated with polynomial delay.*

Proof. Based on Proposition 4, we define M to be the function that maps every $h \in \text{MINTRI}(g)$ to the equivalence class of the proper tree decomposition of h . Next, we now prove that M is as desired.

M has the right range. Let h be a minimal triangulation of g , and let d be a tree decomposition in $M(h)$. Then d is a proper tree decomposition of h , and therefore d is a tree decomposition of g , and we need to show that d is a proper tree decomposition of g . According to Proposition 4, we have $\text{bags}(d) = \text{MAXCLQ}(h)$, and therefore, $\text{saturate}(g, d) = h$. Assume, by way of contradiction, that d is improper. Then d is strictly subsumed by some tree decomposition d' of g , meaning that $d' \sqsubseteq d$ and $\text{bags}(d) \subsetneq \text{bags}(d')$. Let h' be graph $\text{saturate}(g, d')$. From Proposition 3 it follows that h' is a triangulation of g . From $d' \sqsubseteq d$ and the fact that every bag of d is a clique of h , we conclude that $E(h') \subseteq E(h)$. And since h is a minimal triangulation, we get that $E(h') = E(h)$, and therefore $h = h'$. This means that both d and d' are tree decompositions of h , and d is strictly subsumed by d' , which contradicts the fact that d is a proper tree decomposition of h .

M is injective. Let h_1 and h_2 be two minimal triangulations such that $h_1 \neq h_2$. We need to show that $M(h_1) \neq M(h_2)$. Without loss of generality, assume that the edge $\{u, v\}$ is in h_1 but not in h_2 . The nodes u and v are part of some maximal clique of h_1 , so they share a bag in $M(h_1)$. But they are not part of any clique of h_2 , so they do not share any bag in $M(h_2)$. Therefore, $M(h_1) \neq M(h_2)$, as claimed.

M is surjective Given a proper tree decomposition d of g , we need to show that there exists a minimal triangulation h of g such that $d \in M(h)$. Consider the graph $h = \text{saturate}(g, d)$. We will show that h is a minimal triangulation, and that d belongs to $M(h)$.

We first show that h is a minimal triangulation of g . According to Proposition 3, h is a triangulation of g . Assume, by way of contradiction, that h is not minimal. Then there exists a minimal triangulation h' of g that is obtained from h by removing some edges; denote one of these edges by e . Consider a tree decomposition $d' \in M(h')$. The clique containing e in h is not a clique in h' , and therefore $\text{bags}(d) \not\subseteq \text{bags}(d')$. Also note that since $h' \subseteq h$, every maximal clique of h' is contained in some maximal clique of h , and therefore $d' \sqsubseteq d$. Then d' strictly subsumes d , in contradiction to the fact that d is proper.

Finally, we need to show that $\text{bags}(d) = \text{MAXCLQ}(h)$. But this follows immediately from the observation that d is a proper tree decomposition of h , and then applying Proposition 4.

Enumerating proper tree decompositions. Jordan [16] shows that, given a chordal graph h , a tree over the bags that represent the maximal cliques of h is a tree decomposition if and only if it is a maximal spanning tree, where the weight of an edge between two bags is the size of their intersection. Hence, our enumeration problem is reduced to enumerating all maximal spanning trees, which can be solved in polynomial delay [27]. Since Gavril [10] has shown that in chordal graphs the number of maximal cliques of h is at most the number of nodes of h , we have a polynomial delay algorithm for enumerating the tree decompositions. This concludes the proof. \square

3 Enumerating Maximal Independent Sets on Succinct Graph Representations

Our algorithm for enumerating minimal triangulations is done within an abstract framework of *succinct graph representations*, where a graph may be exponentially larger than its representation, and we have access to the nodes and edges through efficient algorithms.

Formally, a *Succinct Graph Representation (SGR)* is a triple (\mathcal{G}, A_V, A_E) , where:

- \mathcal{G} is a function that maps strings x , called *instances*, to graphs $\mathcal{G}(x)$;
- A_V is an enumeration algorithm that, given an instance x , enumerates the nodes of $\mathcal{G}(x)$;
- A_E is an algorithm that, given an instance x and two nodes v and u of $\mathcal{G}(x)$, determines whether v and u are connected by an edge in $\mathcal{G}(x)$.

An SGR (\mathcal{G}, A_V, A_E) is said to be *polynomial* if: (a) A_V enumerates with polynomial delay, and (b) A_E terminates in polynomial time; here, both polynomials are with respect to $|x|$. Observe that in a polynomial SGR, the (representation) size of every node v of $\mathcal{G}(x)$ is polynomial in that of x (since writing v is within the polynomial delay).

Example 1. The *separator graph* of a graph g is the graph that has the set $\text{MINSEP}(g)$ of minimal separators as its node set, and an edge between every two minimal separators that are crossing (i.e., $S, T \in \text{MINSEP}(g)$ such that $S \not\subseteq T$) [22]. Throughout this paper we denote by MSep the SGR $(\mathcal{G}^{\text{ms}}, A_V^{\text{ms}}, A_E^{\text{ms}})$, where:

- \mathcal{G}^{ms} is a function that maps the representation of a graph g to its separator graph $\mathcal{G}^{\text{ms}}(g)$.
- A_V is an enumeration algorithm that, given a graph g , enumerates its set $\text{MINSEP}(g)$ of minimal separators. We can use here the algorithm of Berry et al. [3] that enumerates $\text{MINSEP}(g)$ with polynomial delay. Specifically, the time between two consecutive minimal separators is $O(n^3)$, where n is the number of nodes in g .

- A_E is an algorithm that, given a graph g and two minimal separators S and T , determines whether $S \nmid T$ efficiently (e.g., by removing S and testing whether T is split along multiple connected components).

In particular, **MSep** is a polynomial SGR. \square

A polynomial SGR (\mathcal{G}, A_V, A_E) is said to have a *polynomial expansion of independent sets* if both of the following hold.

1. There is a polynomial p such that for all representations x and independent sets I of $\mathcal{G}(x)$ it holds that $|I| \leq p(|x|)$.
2. There is a polynomial-time algorithm that, given x and an independent set I of $\mathcal{G}(x)$, either determines that I is maximal or returns a node $v \notin I$ such that $I \cup \{v\}$ is independent.

3.1 Enumerating Maximal Independent Sets in SGRs

In this section we prove the following result.

Theorem 5. *Let (\mathcal{G}, A_V, A_E) be a polynomial SGR. If (\mathcal{G}, A_V, A_E) has a polynomial expansion of independent sets then there is an algorithm that, given a representation x , enumerates the maximal independent sets of $\mathcal{G}(x)$ in incremental polynomial time.*

The proof is via the algorithm **EnumMaxIndependent** that is depicted in Figure 3.1. This algorithm is an adaptation of the algorithm for computing full disjunctions in databases [6, 7].

Algorithm description. The algorithm maintains two collections, \mathcal{Q} and \mathcal{O} , for storing intermediate results (which are maximal independent sets). The set \mathcal{O} stores the results that have already been printed, and \mathcal{Q} stores those that are to be printed. Both collections feature logarithmic-time membership-testing and element-removal operations. In addition, the algorithm maintains a collection \mathcal{V} of nodes of $\mathcal{G}(x)$. The collection \mathcal{Q} is initialized with a single result, which is an arbitrary maximal independent set. This result is obtained through the procedure **ExtendInd**(x, I) that extends a given independent set I into a maximal one. (Note that this procedure can be implemented in polynomial time when x has a polynomial expansion of independent sets.) The sets \mathcal{O} and \mathcal{V} are initialized empty.

The algorithm accesses the nodes of $\mathcal{G}(x)$ through an iterator object that is obtained by executing $A_V(x)$, and features two polynomial-time operations:

- Boolean **hasNext()** that determines whether there are additional nodes to enumerate.
- **next()** that returns the next node in the iteration.

Algorithm EnumMaxIndependent(x)

```

1:  $\mathcal{Q} := \{\text{ExtendToMaxInd}(x, \emptyset)\}$ 
2:  $\mathcal{O} := \emptyset$ 
3:  $\mathcal{V} := \emptyset$ 
4:  $\text{iterator} := A_V(x)$ 
5: while  $\mathcal{Q} \neq \emptyset$  do
6:    $I := \mathcal{Q}.\text{pop}()$ 
7:   print  $I$ 
8:    $\mathcal{O}.\text{push}(I)$ 
9:   for all  $v \in \mathcal{V}$  do
10:     $I_v := \{v\} \cup \{u \in \mathcal{O} \mid \neg A_E(x, v, u)\}$ 
11:     $K := \text{ExtendToMaxInd}(x, I_v)$ 
12:    if  $K \notin \mathcal{Q} \cup \mathcal{O}$  then
13:       $\mathcal{Q} := \mathcal{Q} \cup \{K\}$ 
14:   while  $\mathcal{Q} = \emptyset$  and  $\text{iterator.hasNext}()$  do
15:      $\mathcal{V} := \mathcal{V} \cup \{\text{iterator.next}()\}$ 
16:     for all  $v \in \mathcal{V}$  do
17:       for all  $J \in \mathcal{O}$  do
18:          $J_v := \{v\} \cup \{u \in J \mid \neg A_E(x, v, u)\}$ 
19:          $K := \text{ExtendToMaxInd}(x, J_v)$ 
20:         if  $K \notin \mathcal{Q} \cup \mathcal{O}$  then
21:            $\mathcal{Q} := \mathcal{Q} \cup \{K\}$ 

```

Fig. 3. Enumerating maximal independent sets of an input x for a SGR (\mathcal{G}, A_V, A_E)

The algorithm applies the iteration of line 5 until \mathcal{Q} becomes empty, and then it terminates. In every iteration, the algorithm pops an element from \mathcal{Q} , prints it, and stores it in \mathcal{O} (lines 6–8). The algorithm then iterates through the nodes in \mathcal{V} , and for each node v it applies (lines 10–13) what we call *extension of J in the direction of v* : it generates the set J_v that consists of v and all the nodes in J that are non-neighbors of v , and extends J_v into an arbitrary maximal independent set K , again using $\text{ExtendInd}(x, J_v)$. If K is in neither \mathcal{Q} nor \mathcal{O} , then K is added to \mathcal{Q} .

Finally, the algorithm tests whether it is the case that \mathcal{Q} is empty and the node iterator has additional nodes to process (line 14). While this is the case, the algorithm repeats the following procedure (lines 16–21): generate the next node, add it to \mathcal{V} , and extend every result in \mathcal{O} in the direction of every node in \mathcal{V} .

Correctness. We now prove the correctness of the algorithm.

Lemma 1. *Let (\mathcal{G}, A_V, A_E) be a SGR, and let x be an instance. The algorithm `EnumMaxIndependent`(x) enumerates the maximal independent sets of $\mathcal{G}(x)$.*

Proof. Lines 20-21 of the algorithm guarantee that only maximal independent sets are enumerated and that each set is printed only once. We now prove that every maximal independent set of $\mathcal{G}(x)$ is printed.

When the algorithm terminates, $\mathcal{Q} = \emptyset$. Therefore, in the previous iteration, the loop in line 14 could only have terminated due to `iterator.hasNext()` returning false. Therefore, upon termination, we have $\mathcal{V} = \mathcal{V}(\mathcal{G}(g))$.

Suppose, by way of contradiction, that there is some maximal independent set H' that is not printed by the algorithm. Let H be a maximal independent set of $\mathcal{G}(x)$, among all the printed ones, that yields a maximal intersection $H_m = H \cap H'$. Such a set, H , must exist because the algorithm prints at least one maximal independent set. If $H = H'$, then we are done. Otherwise, there is some node $w \in H' \setminus H$.

At this point we have established that both the node w was generated and the set H was printed before the algorithm terminated. We divide into two cases as follows.

1. The set H was printed before the node w was generated. When w is generated (in line 15), then $H \in \mathcal{O}$. During this iteration, the set $H_w = \{w\} \cup \{u \in H \mid \neg A_E(x, w, u)\}$ will be generated and expanded to a maximal independent set $K \supseteq H_w$.
2. The vertex w was generated before the set H was printed. Let us look at the time H is printed and inserted into \mathcal{O} . Since, at this point $w \in \mathcal{V}$, then during the iteration of line 9, the set $H_w = \{w\} \cup \{u \in H \mid \neg A_E(x, w, u)\}$ will be generated and expanded to a maximal independent set $K \supseteq H_w$.

So we have established that before the algorithm terminates, the set $H_w = \{w\} \cup \{u \in H \mid \neg A_E(x, w, u)\}$ will be generated and expanded to a maximal independent set $K \supseteq H_w$. Furthermore, $H_m \cup \{w\} \subseteq H_w \subseteq K$ (because $H_m \subseteq H'$ cannot contain any node that crosses w). According to the algorithm, one of the following options must hold: (1) K is inserted into \mathcal{Q} , (2) K is already in \mathcal{Q} (3) K was in \mathcal{Q} in the past and is now in \mathcal{O} . Since the algorithm prints every maximal independent set that is inserted into \mathcal{Q} , the existence of K contradicts the choice of H and, hence, the existence of H' . \square

Execution time. We now prove that the algorithm enumerates with incremental polynomial time. We denote by $\mathcal{U} = \bigcup_{J \in \mathcal{Q} \cup \mathcal{O}} J$ the set of vertices of \mathcal{G} that are present in one or more maximal independent sets already generated.

Lemma 2. *During the execution of Algorithm `EnumMaxIndependent`, each time line 14 is reached we have that $\mathcal{V} \subseteq \mathcal{U}$.*

Proof. We prove by induction on the number of times line 14 is executed. We denote by \mathcal{V}_j ($j \geq 1$) the set \mathcal{V} when line 14 is reached for the j th time. Clearly, $\mathcal{V}_1 = \emptyset$, so the claim holds. Assume the claim holds for \mathcal{V}_j and we prove for \mathcal{V}_{j+1} .

If $\mathcal{V}_{j+1} = \mathcal{V}_j$ then by the induction hypothesis, the claim holds. Otherwise, $\mathcal{V}_{j+1} = \mathcal{V}_j \cup \{w\}$ where $w \notin \mathcal{V}_j$. If $\mathcal{V}_{j+1} \subseteq \mathcal{U}$ then we are done. Otherwise, we have that $w \notin \mathcal{U}$ at the time the loop in line 16 was executed for the j th time. Let us examine the first iteration of the loop in line 17 that is executed with vertex $v = w$. (Such an iteration must occur after at most $|\mathcal{V}_j|$ iterations of line 16). By this time, either a maximal independent set containing the vertex w was generated and inserted into \mathcal{Q} , hence $w \in \mathcal{U}$. Otherwise, since no maximal independent set in $\mathcal{O} \cup \mathcal{Q}$ contains the node w , then the set (containing w) generated in line 19, will be added to \mathcal{Q} , thus $w \in \mathcal{U}$. Therefore, the next time line 14 is reached, we have $\mathcal{V}_{j+1} = \mathcal{V}_j \cup \{w\} \subseteq \mathcal{U}$, as required. \square

Lemma 3. *Suppose that (\mathcal{G}, A_V, A_E) is an SGR with a polynomial expansion of independent sets, such that $|I| \leq p(x)$ for every instance x and independent set I of $\mathcal{G}(x)$. `EnumMaxIndependent`(x) computes the maximal independent sets of $\mathcal{G}(x)$ with an $O(p(x)^2 N^3(s(x) + \log N) + Np(x)a(x))$ delay, where N is the number of sets already generated, $s(x)$ is a bound on the running time of `ExtendInd`(x, I) and $a(x)$ is the delay of the enumeration algorithm A_V .*

Proof. Since \mathcal{G} has a polynomial expansion of independent sets, then we have that $|\mathcal{U}| \leq p(x)N$. Furthermore, by Lemma 2, we have that $\mathcal{V} \subseteq \mathcal{U}$ each time line 14 is reached. Therefore, the loop in line 16 can run at most $|\mathcal{U}| + 1$ times, or $O(p(x)N)$ times. Hence, the code block of the internal loop in line 17 will be executed a total of $O(p(x)N^2)$ times. After this iteration, one of the following can occur: 1. $\mathcal{Q} \neq \emptyset$, in which case a maximal independent set will be printed 2. `iterator.hasNext()` = false in which case the algorithm will terminate after printing the contents of \mathcal{Q} 3. $\mathcal{Q} = \emptyset$, `iterator.hasNext()` = true and let w denote the node generated in line 15.

Assuming options 1 or 2, we arrive at a delay of $O(p(x)N^2(s(x) + \log N) + a(x))$. We now analyze the runtime for option 3. There are two cases. If $w \notin \mathcal{U}$, then a new maximal independent set containing w will be generated and inserted into \mathcal{Q} (as in case 1 above). Otherwise, the loop of line 16 may be executed without generating a maximal independent set. The number of such loops is confined by the number of vertices that are part of some maximal independent set already generated, and their cardinality is limited by $|\mathcal{U} \setminus \mathcal{V}| = O(Np(x))$. Once these vertices are generated, a node $w \notin \mathcal{U}$ must be generated or the algorithm will terminate. Overall, before the next maximal independent set is generated, at most $O(|\mathcal{U} \setminus \mathcal{V}|) = O(Np(x))$ vertices are returned by the iterator in time $O(Np(x)a(x))$. Also, the number of times the code block in the loop of line 17 is executed is $O(N^3p(x)^2)$.

Summarizing the above three cases, we have that the algorithm will either generate a new maximal independent set or terminate in time $O(p(x)^2 N^3(s(x) + \log N) + Np(x)a(x))$. \square

From Lemmas 1 and 3 we establish Theorem 5.

4 Enumerating Minimal Triangulations

Recall the SGR MSep of Example 1. In this section, we will use known results to reduce the problem of enumerating the minimal triangulations of a graph to the problem of enumerating the maximal independent sets for MSep. We will further show that MSep has polynomial expansion of independent sets. Consequently, we will apply Theorem 5 to conclude that the minimal triangulations can be enumerated in incremental polynomial time.

4.1 Reduction to Enumerating Maximal Sets of Pairwise-Parallel Minimal Separators

We will use the following notation. Let g be a graph. We denote by $\text{CLQMINSEP}(g)$ the set of minimal separators S of g , such that S is a clique of g . Let φ be a subset of $\text{MINSEP}(g)$. We denote by $g_{[\varphi]}$ the graph that results from saturating the minimal separators in φ .

Parra and Scheffler [22] have shown the following connection between minimal triangulations and maximal sets of pairwise-parallel minimal separators (that is, every two are non-crossing).

Theorem 6. (Parra and Scheffler [22]) *Let g be a graph.*

1. *Let $\varphi = \{S_1, \dots, S_k\}$ be a maximal set of pairwise parallel minimal separators of g . Then $h = g_{[\varphi]}$ is a minimal triangulation of g , and $\text{MINSEP}(h) = \varphi$.*
2. *Let h be a minimal triangulation of g . Then $\text{MINSEP}(h)$ is a maximal set of pairwise parallel minimal separators in g , and $h = g_{[\text{MINSEP}(h)]}$.*

We conclude the following corollary, which gives the desired reduction. Recall that the graph $\mathcal{G}^{\text{ms}}(g)$ is defined in Example 1.

Corollary 1. *Given a graph g , there is a polynomial-time-computable bijection between $\text{MINTRI}(g)$ and the maximal independent sets of $\mathcal{G}^{\text{ms}}(g)$.*

4.2 Polynomial Expansion of Independent Sets

It is left to prove that the SGR MSep has polynomial expansion of independent sets. (The definition is in Section 3.) Theorem 6, combined with a result by Rose [24], gives the first of the two conditions.

Corollary 2. *Let g be a graph. If I is a (maximal) independent set in $\mathcal{G}^{\text{ms}}(g)$, then $|I| < V(g)$.*

Proof. Suppose that I is a maximal set of pairwise parallel minimal separators of g . Then by Theorem 6, $h = g_{[I]}$ is a minimal triangulation of g , and $\text{MINSEP}(h) = I$. The graph h is chordal, hence from Rose [24] we get that $|\text{MINSEP}(h)| < |V(h)| = |V(g)|$. \square




Fig. 4. An algorithm for extending a set φ of pairwise-parallel minimal separators

We now turn to proving the second condition. We will do so by describing a general procedure for extending a set of pairwise-parallel minimal separators of a graph g to a maximal such set. Algorithm 4, **ExtendInd**, can apply any known polynomial time triangulation heuristic, referred to as **Triangulate**, as a black box. It uses the following procedures as subroutines.

- **Saturate**(g, S) receives a graph g and a set $S \subseteq V(g)$ of vertices, and saturates S (i.e., modifies g such that S becomes a clique).
- **Triangulate**(g) receives a graph g and returns a triangulation g' of g . We assume that this procedure runs in time that is a polynomial function of $(|V(g)| + |E(g)|)$.
- **MinTriSandwich**(g, g') receives a graph g and a triangulation g' of g , and return *minimal* triangulation of g . We note that, using one of known algorithms [5, 8, 23], this procedure runs in time that is polynomial in the size of the graph.
- **ExtractMinSeps**(h) receives a chordal graph h and returns its set of minimal separators. Using the algorithm of Kumar [20], the execution time of this procedure is linear in h .

Algorithm 4 receives as input a graph g and a set φ of pairwise-parallel minimal separators. It then proceeds by saturating the separators in φ , resulting in $g_{[\varphi]}$. At this stage it will pass $g_{[\varphi]}$ to the triangulation heuristic **Triangulate**. We note that **Triangulate** does not have to result in a minimal triangulation. It can involve, for example, a procedure which constructs a tree decomposition, from which a triangulation can be extracted (Proposition 3).

Transforming a non-minimal triangulation to one that is minimal, by removing *fill* edges, is called the *minimal triangulation sandwich problem* [13]. Various polynomial-time algorithms for this problem exist for this problem [8, 23], and these were reported to perform well in practice [5].

So, at this stage we have a minimal triangulation g_t of $g_{[\varphi]}$. Theorem 7 shows that g_t is also a minimal triangulation of g . Lemma 4 shows that the set of minimal separators of g_t contains φ . Finally, we can apply the algorithm of Kumar [20] to extract the minimal separators of the (chordal) graph g_t , in linear time.

Correctness. To prove correctness of the algorithm, we need the following result by Heggernes [13].

Theorem 7. (Heggernes [13]) *Given a graph g , let φ be an arbitrary set of pairwise non-crossing minimal separators of g . Obtain a graph $g_{[\varphi]}$ by saturating each separator in φ .*

1. $\varphi \subseteq \text{CLQMINSEP}(g_{[\varphi]})$, that is, φ consists of clique minimal separators of $g_{[\varphi]}$.
2. $\text{CLQMINSEP}(g) \subseteq \text{MINSEP}(g_{[\varphi]})$; that is, every clique minimal separator of g is a (clique) minimal separator of $g_{[\varphi]}$.
3. Every minimal triangulation of $g_{[\varphi]}$ is a minimal triangulation of g .

The next Lemma 4 builds on Theorems 6 and 7.

Lemma 4. *Let g be a graph, and φ a set of pairwise-parallel minimal separators of g . Let g_t a minimal triangulation of $g_{[\varphi]}$. Then $\varphi \subseteq \text{MINSEP}(g_t)$.*

Proof. By Part 1 of Theorem 7 we have that $\varphi \subseteq \text{CLQMINSEP}(g_{[\varphi]})$. Since g_t is a minimal triangulation of $g_{[\varphi]}$ then by Theorem 6, g_t is the result of saturating a maximal set, say φ' , of pairwise-parallel minimal separators of $g_{[\varphi]}$. Therefore, by Part 2 of Theorem 7 we have $\text{CLQMINSEP}(g_{[\varphi]}) \subseteq \text{MINSEP}(g_t)$. This implies that $\varphi \subseteq \text{MINSEP}(g_t)$, as claimed.

We then conclude the correctness of the algorithm.

Lemma 5. *Let φ be a set of pairwise-parallel minimal separators of a graph g . Algorithm 4 returns a maximal set I of minimal separators of g such that $\varphi \subseteq I$. Furthermore, the algorithm terminates in polynomial time.*

Proof. Assuming the correctness of the procedures **Triangulate**, and **MinTriSandwich**, the graph g_t is a minimal triangulation of $g_{[\varphi]}$. By Part 3 of Theorem 7, we have that g_t is a minimal triangulation of g . Consequently, from Theorem 6 we get that $\text{MINSEP}(g_t) = I$ is a maximal set of non-crossing minimal separators of g . By Lemma 4 it holds that $\varphi \subseteq \text{MINSEP}(g_t)$, making I an extension of φ .

All of the procedures in Algorithm 4 run in time that is polynomial in the size of the graph making it polynomial as well.

From Corollary 2 and Lemma 5 we get the main result of this part.

Theorem 8. *The SGR MSEP has a polynomial expansion of independent sets.*

4.3 Main Result

From Theorems 8 and 5 we conclude that it is possible to enumerate the maximal independent sets of **MSEP** in incremental polynomial time. Applying the bijections of Theorems 1 and 4, we get the main result of this paper.

Theorem 9. *There are algorithms that, given a graph g , enumerate in incremental polynomial time:*

1. *The minimal triangulations of g .*
2. *The proper tree decompositions of g .*

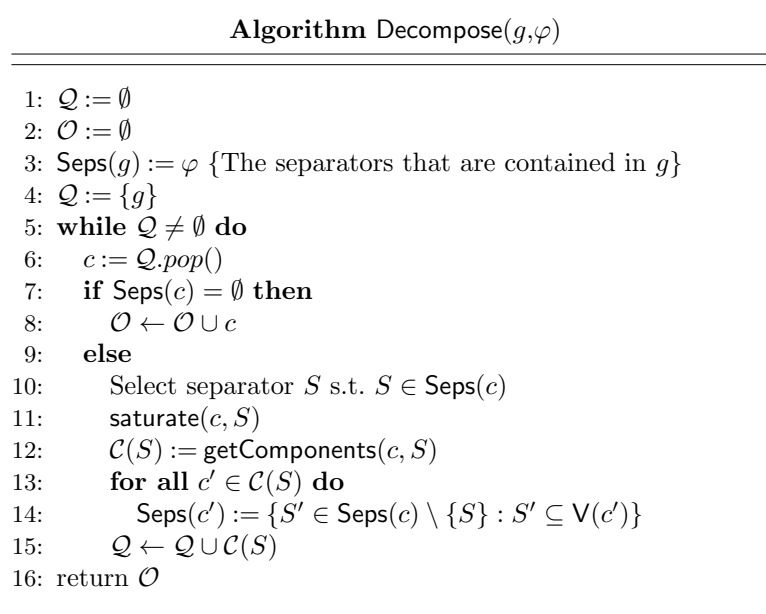


Fig. 5. An algorithm that decomposes a connected graph using a set of pairwise non-crossing minimal separators, φ

5 A Generic Algorithm for Expanding Independent Sets

In this section we provide a generic procedure for extending a parallel set of minimal separators to one that is maximal. The procedure is provided in addition the one described in section 4 and enables triangulating the graph by applying separator-based approaches [1].

Algorithm 6 is a procedure for extending a set, φ , of pairwise parallel minimal separators of a graph g , into one that is maximal. Let c be a connected component of g . We denote by $\text{SEPS}(c) = \{S \in \varphi : S \subseteq V(c)\}$. Clearly, $\text{SEPS}(g) = \varphi$.

The algorithm uses the following subroutines:

- **saturate**(c, S) receives a graph c and a set of vertices $S \subseteq V(c)$ and modifies c such that S is a clique in c . The complexity of this procedure is $O(n^2)$.
- **getComponents**(g, S) receives a graph g and a set of pairwise connected vertices $S \subseteq V(g)$ and returns a set of induced subgraphs of g as follows. Let c_1, c_2, \dots, c_k be the set of connected components of $g_{|V(g) \setminus S}$. The procedure returns the set $\mathcal{C}(S) = \bigcup_{i \in [1, k]} \{g_{|c_i \cup N_g(c_i)}\}$. Computing the connected components of a graph can be performed in linear time using depth-first search.
- **findMinSep**($\{u, v\}, c$) receives a graph c with two non-adjacent vertices $u, v \in c$ and returns a minimal separator, S , for this pair. A separator S that is “close” to u , i.e., $S \subseteq N_c(u)$, can be found in linear time.

Algorithm ExtendInd(g, φ)

```

1:  $\mathcal{Q} := \emptyset$ 
2:  $I := \varphi$ 
3: if  $\varphi \neq \emptyset$  then
4:    $\mathcal{Q} := \text{Decompose}(g, \varphi)$  {Call procedure 5}
5: else
6:    $\mathcal{Q} := \{g\}$ 
7: while  $\mathcal{Q} \neq \emptyset$  do
8:    $c := \mathcal{Q}.\text{pop}()$ 
9:   if !isClique( $c$ ) then
10:     $(u, v) := \{u, v \in V(c) : (u, v) \notin E(c)\}$ 
11:     $S := \text{findMinSep}(\{u, v\}, c)$ 
12:    saturate( $c, S$ )
13:     $\mathcal{C}(S) := \text{getComponents}(c, S)$  {Each  $c' \in \mathcal{C}(S)$  contains  $S \cap N_g(c')$  as a clique}
14:    for all  $c' \in \mathcal{C}(S)$  do
15:       $I := I \cup \{N_g(c') \cap S\}$  {In order to include contained separators}
16:     $\mathcal{Q} := \mathcal{Q} \cup \mathcal{C}(S)$ 
17: Return  $I$ 

```

Fig. 6. An algorithm that finds a maximal set, I , of pairwise non-crossing minimal separators of a graph g that contains the input set φ

- isClique(c) receives a graph c , and returns true if c is a clique. The complexity of this procedure is $O(n^2)$.
- Decompose (Algorithm 5) receives a connected graph, g , and a set of pairwise parallel separators, φ , and returns the set of connected components that result from decomposing g according to the separators of φ .

Algorithm 6 first decomposes the graph into connected components, according to the separator set φ , received as input. Then, each connected component, c , is processed, in turn. If it has some pair, (u, v) , of unconnected vertices, then a minimal separator for them, S , is found, saturated and used to further decompose the connected component. The loop in line 14 iterates over the resulting connected components and updates I with the set of minimal separators included in S .

Lemmas 6 and 7 prove general properties of the decomposition process by parallel minimal separators and are used to prove the correctness and complexity of Algorithms 5 (Decompose) and 6 (ExtendInd).

Lemma 6. *Let g be a connected graph, φ be a set of pairwise non-crossing minimal separators and $\mathcal{C}(S) = \text{getComponents}(g, S)$ where $S \in \varphi$ and the set of nodes in S are pairwise connected. Then:*

1. S , and its subsets, are no longer minimal separators in any of the subgraphs in $\mathcal{C}(S)$.
2. For all $S' \in \varphi \setminus \{S\}$, there exists some subgraph $c \in \mathcal{C}(S)$ such that $S' \subseteq c$.

Proof. The subroutine `getComponents` returns the set of subgraphs $\mathcal{C}(S) = \bigcup_i g_i$ where $g_i = g_{|c_i \cup N_g(c_i)}$ is the subgraph induced by the connected component $c_i \in g \setminus S$ and its neighbors, $N_g(c_i)$. Since S is a minimal separator, then for every $c_i \in g \setminus S$ we have that $N_g(c_i) \subseteq S$.

Let us assume, by contradiction, that there is some $g_i \in \mathcal{C}(S)$, with two non-adjacent nodes $u, v \in V(g_i)$ that are separated by $S' \subseteq S$. It cannot be the case that $u, v \in N_g(c_i)$ because $N_g(c_i) \subseteq S$ and S is pairwise connected. If $u, v \in V(c_i)$ then it must be that $u, v \in V(g) \setminus S$ and thus should have been in different components of $\mathcal{C}(S)$. Finally, assume that $u \in N_g(c_i) \setminus S'$ and $v \in V(c_i)$. Since c_i is a connected component and u is adjacent to some node $w \in c_i$, we have a path from u to v (via w) that does not pass through S' , in contradiction to the fact that S' is a minimal (u, v) -separator.

If we assume that there is an $S' \in \varphi$ that is not contained in any component of $\mathcal{C}(S)$, then S' must span at least two components in $\mathcal{C}(S)$. This means that S crosses S' , and we arrive at a contradiction that the members of φ are pairwise non-crossing. \square

Lemma 7. *Let g be a connected graph, S be a clique minimal separator of g and $\mathcal{C}(S) = \text{getComponents}(g, S)$. Let $c_1, c_2 \in \mathcal{C}(S)$, then $\text{SEPS}(c_1) \cap \text{SEPS}(c_2) = \emptyset$, where $\text{SEPS}(c)$ denotes the minimal separators in component c .*

Proof. Let us assume, by contradiction, that there is some $S' \in \text{SEPS}(c_1) \cap \text{SEPS}(c_2)$. By Lemma 6, we know that $S' \not\subseteq S$. Therefore, there is some node $w \in S' \setminus S$ such that $w \in c_1 \cap c_2$. But this means that c_1 and c_2 are connected in $g_{|V(g) \setminus S}$ (via w) in contradiction to the assumption that these are distinct components of $\mathcal{C}(S)$. \square

Lemma 8 proves the correctness of Algorithm 5, while Lemma 10 applies the proof of Lemma 9 to prove its complexity.

Lemma 8. *Let g denote a connected graph provided as input to `Decompose`. For every pair of nodes $(u, v) \in V(g)$, u and v will reside in distinct components $c_u, c_v \in \mathcal{O}$ if and only if they are separated by some $S \in \varphi$.*

Proof. If u and v are in distinct components $c_u, c_v \in \mathcal{O}$, then this must be the result of a separation applied in line 12 using one of the separators of $S \in \varphi$.

Let $S \in \varphi$ be a (u, v) -separator and assume, by contradiction, that at the end of the procedure u and v reside in a common component $c \in \mathcal{O}$. We first show that (u, v) cannot be connected by an edge in c . If this were the case, then either $(u, v) \in E(g)$, contradicting the fact that they are even separable by S . Otherwise, the edge is a result of saturating some separator $S' \in \varphi$, in which case we get that S crosses S' , a contradiction. Therefore, there is a path from u to v in c , and thus, $S \cap c \neq \emptyset$. If $S \not\subseteq c$, then it must be crossed by some other separator S' , and we arrive at a contradiction. Otherwise, $S \subseteq c$, and by

definition $S \in \text{SEPS}(c)$. But this cannot be the case because c is inserted into \mathcal{O} only if $\text{SEPS}(c) = \emptyset$ (line 7). Hence, we arrive at a contradiction for the existence of a component c containing both u and v . \square

The lemmas that follow prove the correctness and complexity of Algorithm 6.

Lemma 9. *Every separator discovered in line 11 is generated and added to the resulting set, I , exactly once.*

Proof. Assume, by contradiction, that there is some minimal separator S that is generated twice in line 11. First from component c and then from component c' . There are two options: (1) c' is contained in some component of $\text{getComponents}(c, S)$. But this contradicts Lemma 6 which states that S cannot be a separator in any of the subgraphs of $\text{getComponents}(c, S)$. (2) Otherwise, let S' be the earliest separator such that the members of c and c' belong to distinct components (such a separator must exist because we start out with a connected graph g). But this contradicts Lemma 7 which states that the separators of c and c' must be disjoint. \square

Lemma 10. *Algorithm Decompose runs in polynomial time.*

Proof. Using the same arguments as those in the proof of Lemma 9, it is shown that every separator $S \in \varphi$ is processed exactly once by Algorithm 5. This, combined with the fact that the subroutines called by Decompose run in polynomial time, brings us to the desired result. \square

Lemma 11. *The set of minimal vertex separators in I are pairwise non-crossing.*

Proof. We show by induction on the number of iterations of the loop in line 7 that the set of separators in I are pairwise non-crossing in g , and that they represent cliques in all the components in \mathcal{Q} that contain them.

Since the input φ is a set of pairwise non-crossing minimal separators, that undergo saturation, the claim holds before the loop in line 7. Assume the claim holds until some iteration $j > 0$. Let c be the component processed in iteration $j+1$. If the minimal separator S (found in line 11) crosses some separator $S' \in I$, then it means that there exist $x, y \in S'$ such that $x, y \in c$ and $(x, y) \notin E(c)$ in contradiction to the induction assumption stating that S' is a clique in c . \square

Lemma 12. *The set of minimal vertex separators I , returned by algorithm 6, is a maximal set of pairwise non-crossing minimal separators of g .*

Proof. We have already shown in Lemma 11 that the set of separators in I is pairwise non-crossing, we now show that it is maximal.

Let S' be a minimal separator of graph g such that for every $S \in I$, S and S' are non-crossing. We show that $S' \in I$. Let c be the *latest* component, processed

by algorithm 6, such that $S' \subseteq c$. Such a component must exist because S' is not crossed by any separator in I (Lemma 6). Let S be the minimal separator generated in line 11. If $S = S'$, then we are done because $S \in I$. Otherwise, by definition of non-crossing, there exists some component $c' \in \mathcal{C}(S)$ such that $S' \subseteq c' \subseteq c$. But this is a contradiction because c is the latest component processed that contains S' . \square

Theorem 10. *Algorithm 6 runs in time that is polynomial in $|V(g)| = n$.*

Proof. Lemma 12 establishes that the set of separators, I , returned by algorithm 6, is maximal pairwise-parallel. By corollary 2, $|I| < n$. By lemma 9, each member of I is generated and inserted into I exactly once. Therefore, line 15, as well as the number of iterations of the loop in line 7, in algorithm 6, will be executed at most n times.

Since all of the subroutines referred to in algorithm 6 can be performed in polynomial time, the complexity of algorithm 6 is polynomial in n . \square

Theorem 11. $\text{MSep} = (\mathcal{G}, A_V, A_E)$ *has a polynomial expansion of independent sets.*

Proof. By corollary 2 the size of each maximal independent set, I , of $\mathcal{G}(g)$ is at most $|V(g)| = n$, thereby satisfying the first requirement of polynomial expansion.

By Lemma 12 and Theorem 10, we can check that I is a maximal set of pairwise parallel separators, in polynomial time, by verifying that the output of algorithm 6, when provided with input I , is simply I . \square

Hence, as a corollary we get the main result of this paper.

Corollary 3. *The minimal triangulations of a given graph can be enumerated in incremental polynomial time. Hence, due to Theorem 4, the proper tree decompositions of a given graph can be enumerated in incremental polynomial time.*

References

1. E. Amir. Approximation algorithms for treewidth. *Algorithmica*, 56(4):448–479, 2010.
2. S. Arnborg, D. G. Corneil, and A. Proskurowski. Complexity of finding embeddings in ak-tree. *SIAM Journal on Algebraic Discrete Methods*, 8(2):277–284, 1987.
3. A. Berry, J. P. Bordat, and O. Cogis. Generating all the minimal separators of a graph. In P. Widmayer, G. Neyer, and S. Eidenbenz, editors, *WG*, volume 1665 of *Lecture Notes in Computer Science*, pages 167–172. Springer, 1999.
4. A. Berry, P. Heggernes, and G. Simonet. *Graph-Theoretic Concepts in Computer Science: 29th International Workshop, WG 2003. Elspeet, The Netherlands, June 19-21, 2003. Revised Papers*, chapter The Minimum Degree Heuristic and the Minimal Triangulation Process, pages 58–70. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.
5. J. R. Blair, P. Heggernes, and J. A. Telle. A practical algorithm for making filled graphs minimal. *Theoretical Computer Science*, 250(1–2):125 – 141, 2001.

6. S. Cohen, I. Fadida, Y. Kanza, B. Kimelfeld, and Y. Sagiv. Full disjunctions: Polynomial-delay iterators in action. In *Proceedings of the 32nd International Conference on Very Large Data Bases, Seoul, Korea, September 12-15, 2006*, pages 739–750. ACM, 2006.
7. S. Cohen, B. Kimelfeld, and Y. Sagiv. Generating all maximal induced subgraphs for hereditary and connected-hereditary graph properties. *J. Comput. Syst. Sci.*, 74(7):1147–1159, 2008.
8. D. Dahlhaus. *Graph-Theoretic Concepts in Computer Science: 23rd International Workshop, WG'97 Berlin, Germany, June 18–20, 1997 Proceedings*, chapter Minimal elimination ordering inside a given chordal graph, pages 132–143. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997.
9. G. A. Dirac. On rigid circuit graphs. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 25, Universität Hamburg, 1961.
10. F. Gavril. The intersection graphs of subtrees in trees are exactly the chordal graphs. *J. Combinatorial Theory*, 16:47–56, 1974.
11. G. Gottlob, G. Greco, and F. Scarcello. Pure nash equilibria: Hard and easy games. *J. Artif. Intell. Res. (JAIR)*, 24:357–406, 2005.
12. G. Gottlob, M. Grohe, N. Musliu, M. Samer, and F. Scarcello. Hypertree decompositions: Structure, algorithms, and applications. In *WG*, volume 3787 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2005.
13. P. Heggernes. Minimal triangulations of graphs: A survey. *Discrete Mathematics*, 306(3):297 – 317, 2006. Minimal Separation and Minimal Triangulation.
14. P. Heggernes. Treewidth, partial k-trees, and chordal graphs. *unpublished*, 2006.
15. D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. On generating all maximal independent sets. *Inf. Process. Lett.*, 27(3):119–123, 1988.
16. M. Jordan. *An Introduction to Probabilistic Graphical Models*, chapter 17. University of California, Berkeley, 2002.
17. B. Kenig and A. Gal. On the impact of junction-tree topology on weighted model counting. In *SUM*, volume 9310 of *Lecture Notes in Computer Science*, pages 83–98. Springer, 2015.
18. T. Kloks, D. Kratsch, and J. P. Spinrad. On treewidth and minimum fill-in of asteroidal triple-free graphs. *Theor. Comput. Sci.*, 175(2):309–335, 1997.
19. P. G. Kolaitis and M. Y. Vardi. Conjunctive-query containment and constraint satisfaction. *J. Comput. Syst. Sci.*, 61(2):302–332, 2000.
20. P. Kumar and C. Madhavan. Minimal vertex separators of chordal graphs. *Discrete Applied Mathematics*, 89(1–3):155 – 168, 1998.
21. S. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society, B*, 50(2):157–224, 1988.
22. A. Parra and P. Scheffler. Characterizations and algorithmic applications of chordal graph embeddings. *Discrete Applied Mathematics*, 79(1-3):171–188, 1997.
23. B. W. Peyton. Minimal orderings revisited. *SIAM J. Matrix Analysis Applications*, 23(1):271–294, 2001.
24. D. J. Rose. Triangulated graphs and the elimination process. *Journal of Mathematical Analysis and Applications*, 32(3):597 – 609, 1970.
25. R. E. Tarjan and M. Yannakakis. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM J. Comput.*, 13(3):566–579, July 1984.
26. S. Tu and C. Ré. DuncCap: Query plans using generalized hypertree decompositions. In *SIGMOD*, pages 2077–2078. ACM, 2015.

27. T. Yamada, S. Kataoka, and K. Watanabe. Listing all the minimum spanning trees in an undirected graph. *Int. J. Comput. Math.*, 87(14):3175–3185, 2010.